

-1-

Date: 4-12-00 Express Mail Label No. EL 192655203US

Inventors: Thomas Mark Levergood, Lawrence C. Stewart, Stephen
Jeffrey Morris, Andrew C. Payne, and George Winfield
Treese

Attorney's Docket No.: 1984.1001-004

WEB ADVERTISING METHOD

~~INTERNET SERVER ACCESS CONTROL AND MONITORING SYSTEMS~~

RELATED APPLICATION

This application is a Divisional of U.S. of Application No. 09/005,479 filed
January 12, 1998, which is a Continuation of U.S. Application No. 08/474,096, filed
5 June 7, 1995, the entire teachings of which are incorporated herein by reference.

BACKGROUND TO THE INVENTION

The Internet, which started in the late 1960s, is a vast computer network
consisting of many smaller networks that span the entire globe. The Internet has grown
exponentially, and millions of users ranging from individuals to corporations now use
10 permanent and dial-up connections to use the Internet on a daily basis worldwide. The
computers or networks of computers connected within the Internet, known as "hosts",
allow public access to databases featuring information in nearly every field of expertise
and are supported by entities ranging from universities and government to many
commercial organizations.

15 The information on the Internet is made available to the public through
"servers". A server is a system running on an Internet host for making available files or
documents contained within that host. Such files are typically stored on magnetic
storage devices, such as tape drives or fixed disks, local to the host. An Internet server
may distribute information to any computer that requests the files on a host. The

002140-5623750

[illegible]

5

15

25

primarily over TCP/IP and uses the standard Internet setup, where a server issues the data and a client displays or processes it. One format for information transfer is to create documents using Hypertext Markup Language (HTML). HTML pages are made up of standard text as well as formatting codes which indicate how the page should be displayed. The Web client, a browser, reads these codes in order to display the page. The hypertext conventions and related functions of the world wide web are described in the appendices of U.S. Patent Application Serial No. 08/328,133, filed on October 24, 1994, by Payne et al. which is incorporated herein by reference.

Each Web page may contain pictures and sounds in addition to text. Hidden behind certain text, pictures or sounds are connections, known as "hypertext links" ("links"), to other pages within the same server or even on other computers within the Internet. For example, links may be visually displayed as words or phrases that may be underlined or displayed in a second color. Each link is directed to a web page by using a special name called a URL (Uniform Resource Locator). URLs enable a Web browser to go directly to any file held on any Web server. A user may also specify a known URL by writing it directly into the command line on a Web page to jump to another Web page.

The URL naming system consists of three parts: the transfer format, the host name of the machine that holds the file, and the path to the file. An example of a URL may be:

http://www.college.univ.edu/Adir/Bdir/Cdir/page.html,

where "http" represents the transfer protocol; a colon and two forward slashes (://) are used to separate the transfer format from the host name; "www.college.univ.edu" is the host name in which "www" denotes that the file being requested is a Web page; "/Adir/Bdir/Cdir" is a set of directory names in a tree structure, or a path, on the host

machine; and "page.html" is the file name with an indication that the file is written in HTML.

The Internet maintains an open structure in which exchanges of information are made cost-free without restriction. The free access format inherent to the Internet, however, presents difficulties for those information providers requiring control over their Internet servers. Consider for example, a research organization that may want to make certain technical information available on its Internet server to a large group of colleagues around the globe, but the information must be kept confidential. Without means for identifying each client, the organization would not be able to provide information on the network on a confidential or preferential basis. In another situation, a company may want to provide highly specific service tips over its Internet server only to customers having service contracts or accounts.

Access control by an Internet server is difficult for at least two reasons. First, when a client sends a request for a file on a remote Internet server, that message is routed or relayed by a web of computers connected through the Internet until it reaches its destination host. The client does not necessarily know how its message reaches the server. At the same time, the server makes responses without ever knowing exactly who the client is or what its IP address is. While the server may be programmed to trace its clients, the task of tracing is often difficult, if not impossible. Secondly, to prevent unwanted intrusion into private local area networks (LAN), system administrators implement various data-flow control mechanisms, such as the Internet "firewalls", within their networks. An Internet firewall allows a user to reach the Internet anonymously while preventing intruders of the outside world from accessing the user's LAN.

SUMMARY OF THE INVENTION

The present invention relates to methods of processing service requests from a client to a server through a network. In particular the present invention is applicable to processing client requests in an HTTP (Hypertext Transfer Protocol) environment, such as the World-Wide Web (Web). One aspect of the invention involves forwarding a service request from the client to the server and appending a session identification (SID) to the request and to subsequent service requests from the client to the server within a session of requests. In a preferred embodiment, the present method involves returning the SID from the server to the client upon an initial service request made by the client.

10 A valid SID may include an authorization identifier to allow a user to access controlled files.

In a preferred embodiment, a client request is made with a Uniform Resource Locator (URL) from a Web browser. Where a client request is directed to a controlled file without an SID, the Internet server subjects the client to an authorization routine prior to issuing the SID, the SID being protected from forgery. A content server initiates the authorization routine by redirecting the client's request to an authentication server which may be at a different host. Upon receiving a redirected request, the authentication server returns a response to interrogate the client and then issues an SID to a qualified client. For a new client, the authentication server may open a new account and issue an SID thereafter. A valid SID typically comprises a user identifier, an accessible domain, a key identifier, an expiration time such as date, the IP address of the user computer, and an unforgeable digital signature such as a cryptographic hash of all of the other items in the SID encrypted with a secret key. The authentication server then forwards a new request consisting of the original URL appended by the SID to the client in a REDIRECT. The modified request formed by a new URL is automatically forwarded by the client browser to the content server.

15

20

25

When the content server receives a URL request accompanied by an SID, it logs the URL with the SID and the user IP address in a transaction log and proceeds to validate the SID. When the SID is so validated, the content server sends the requested document for display by the client's Web browser.

- 5 In the preferred embodiment, a valid SID allows the client to access all controlled files within a protection domain without requiring further authorization. A protection domain is defined by the service provider and is a collection of controlled files of common protection within one or more servers.

- 10 When a client accesses a controlled Web page with a valid SID, the user viewing the page may want to traverse a link to view another Web page. There are several possibilities. The user may traverse a link to another page in the same path. This is called a "relative link". A relative link may be made either within the same domain or to a different domain. The browser on the client computer executes a relative link by rewriting the current URL to replace the old controlled page name with a new one. The
15 new URL retains all portions of the old, including the SID, except for the new page name. If the relative link points to a page in the same protection domain, the SID remains valid, and the request is honored. However, if the relative link points to a controlled page in a different protection domain, the SID is no longer valid, and the client is automatically redirected to forward the rewritten URL to the authentication
20 server to update the SID. The updated or new SID provides access to the new domain if the user is qualified.

- 25 The user may also elect to traverse a link to a document in a different path. This is called an "absolute link". In generating a new absolute link, the SID is overwritten by the browser. In the preferred embodiment, the content server, in each serving of a controlled Web page within the domain, filters the page to include the current SID in each absolute URL on the page. Hence, when the user elects to traverse an absolute

link, the browser is facilitated with an authenticated URL which is directed with its SID to a page in a different path. In another embodiment, the content server may forego the filtering procedure as above-described and redirect an absolute URL to the authentication server for an update.

- 5 An absolute link may also be directed to a controlled file in a different domain. Again, such a request is redirected to the authentication server for processing of a new SID. An absolute link directed to an uncontrolled file is accorded an immediate access.

- 10 In another embodiment, a server access control may be maintained by programming the client browser to store an SID or a similar tag for use in each URL call to that particular server. This embodiment, however, requires a special browser which can handle such communications and was generally not suitable for early browser formats common to the Web. However, it may now be implemented in cookie compatible browsers.

- 15 Another aspect of the invention is to monitor the frequency and duration of access to various pages both controlled and uncontrolled. A transaction log within a content server keeps a history of each client access to a page including the link sequence through which the page was accessed. Additionally, the content server may count the client requests exclusive of repeated requests from a common client. Such records provide important marketing feedback including user demand, access pattern, and
20 relationships between customer demographics and accessed pages and access patterns.

Fig. 1

- 25 The above and other features of the invention including various novel details of construction and combinations of parts will now be more particularly described with reference to the accompanying drawings and pointed out in the claims. It will be understood that the particular devices and methods embodying the invention are shown by way of illustration only and not as limitations of the invention. The principles and

features of this invention may be employed in varied and numerous embodiments without departing from the scope of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a diagram illustrating the Internet operation.

5 Figure 2A is a flowchart describing the preferred method of Internet server access control and monitoring.

Figure 2B is a related flowchart describing the details of the authentication process.

10 Figure 3 illustrates an example of a client-server exchange session involving the access control and monitoring method of the present invention.

Figure 4 is an example of a World Wide Web page.

Figure 5 is an example of an authorization form page.

Figure 6 is a diagram describing the details of the translation of telephone numbers to URLs.

15 DETAILED DESCRIPTION OF THE INVENTION

Referring now to the drawings, Figure 1 is a graphical illustration of the Internet. The Internet 10 is a network of millions of interconnected computers 12 including systems owned by Internet providers 16 and information systems (BBS) 20 such as CompuServe or America Online. Individual or corporate users may establish
20 connections to the Internet in several ways. A user on a home PC 14 may purchase an account through the Internet provider 16. Using a modem 22, the PC user can dial up the Internet provider to connect to a high speed modem 24 which, in turn, provides a full service connection to the Internet. A user 18 may also make a somewhat limited
25 connection to the Internet through a BBS 20 that provides an Internet gateway connection to its customers.

Figure 2A is a flowchart detailing the preferred process of the present invention and Figure 4 illustrates a sample Web page displayed at a client by a browser. The page includes text 404 which includes underlined link text 412. The title bar 408 and URL bar 402 display the title and URL of the current web page, respectively. As shown in

5 Figure 4, the title of the page is "Content Home Page" and the corresponding URL is "http://content.com/homepage". When a cursor 414 is positioned over link text 412b, the page which would be retrieved by clicking a mouse is typically identified in a status bar 406 which shows the URL for that link. In this example the status bar 406 shows that the URL for the pointed link 412b is directed to a page called "advertisement" in a

10 commercial content server called "content". By clicking on the link text, the user causes the browser to generate a URL GET request at 100 in Figure 2A. The browser forwards the request to a content server 120, which processes the request by first determining whether the requested page is a controlled document 102. If the request is directed to an uncontrolled page, as in "advertisement" page in this example, the content server

15 records the URL and the IP address, to the extent it is available, in the transaction log 114. The content server then sends the requested page to the browser 116 for display on the user computer 117.

If the request is directed to a controlled page, the content server determines whether the URL contains an SID 102. For example, a URL may be directed to a

20 controlled page name "report", such as "http://content.com/report", that requires an SID. If no SID is present, as in this example, the content server sends a "REDIRECT" response 122 to the browser 100 to redirect the user's initial request to an authentication server 200 to obtain a valid SID. The details of the authentication process are described in Figure 2B and will be discussed later, but the result of the process is an SID provided

25 from the authentication server to the client. In the above example, a modified URL appended with an SID may be: "http://content.com/[SID]/report". The preferred SID is a sixteen character ASCII string that encodes 96 bits of SID data, 6 bits per character. It contains a 32-bit digital signature, a 16-bit expiration date with a granularity of one

hour, a 2-bit key identifier used for key management, an 8-bit domain comprising a set of information files to which the current SID authorizes access, and a 22-bit user identifier. The remaining bits are reserved for expansion. The digital signature is a cryptographic hash of the remaining items in the SID and the authorized IP address
5 which are encrypted with a secret key which is shared by the authentication and content servers.

If the initial GET URL contains a SID, the content server determines whether the request is directed to a page within the current domain 106. If the request having a SID is directed to a controlled page of a different domain, the SID is no longer valid
10 and, again, the user is redirected to the authentication server 122.

If the request is for a controlled page within the current domain, the content server proceeds to log the request URL, tagged with SID, and the user IP address in the transaction log 108. The content server then validates the SID 110. Such validation includes the following list of checks: (1) the SID's digital signature is compared against
15 the digital signature computed from the remaining items in the SID and the user IP address using the secret key shared by the authentication and content servers; (2) the domain field of the SID is checked to verify that it is within the domain authorized; and (3) the EXP field of the SID is checked to verify that it is later than the current time.

If the validation passes, the content server searches the page to be forwarded for
20 any absolute URL links contained therein 112, that is, any links directed to controlled documents in different content servers. The content server augments each absolute URL with the current SID to facilitate authenticated accesses across multiple content servers. The requested page as processed is then transmitted to the client browser for display 117. The user viewing the requested Web page may elect to traverse any link on
25 that page to trigger the entire sequence again 100.

Figure 2B describes the details of the authentication process. The content server may redirect the client to an authentication server. The REDIRECT URL might be:

"http://auth.com/authenticate?domain=[domain]&URL=http://

content.com/report". That URL requests authentication and specifies the domain and

- 5 the initial URL. In response to the REDIRECT, the client browser automatically sends a GET request with the provided URL.

Whenever the content server redirects the client to the authentication server 200, the authentication server initiates the authorization process by validating that it is for an approved content server and determining the level of authentication required for the

- 10 access requested 210. Depending on this level, the server may challenge the user 212 for credentials. If the request is for a low level document, the authentication may issue an appropriate SID immediately 228 and forego the credential check procedures. If the document requires credentials, the authentication server sends a "CHALLENGE" response which causes the client browser to prompt the user for credentials 214. A
15 preferred credential query typically consists of a request for user name and password. If the user is unable to provide a password, the access is denied. The browser forms an authorization header 300 from the information provided, and resends a GET request to the authentication server using the last URL along with an authorization header. For example, a URL of such a GET request may be:

- 20 "http://auth.com/authenticate?domain=[domain]&URL=http://
content.com/report and the authorization header may
be:"AUTHORIZE:[authorization]".

- Upon receiving the GET request, the authentication server queries an account database 216 to determine whether the user is authorized 218 to access the requested
25 document. A preferred account database may contain a user profile which includes information for identifying purposes, such as client IP address and password, as well as user demographic information, such as user age, home address, hobby, or occupation,

for later use by the content server. If the user is authorized, an SID is generated 228 as previously described. If the user is not cleared for authorization, the authentication server checks to see if the user qualifies for a new account 220. If the user is not qualified to open a new account, a page denying access 222 is transmitted to the client browser 100. If the user is qualified, the new user is sent a form page such as illustrated in Figure 5 to initiate a real-time on-line registration 224. The form may, for example, require personal information and credit references from the user. The browser is able to transmit the data entered by the user in the blanks 502 as a "POST" message to the authentication server. A POST message causes form contents to be sent to the server in a data body other than as part of the URL. If the registration form filled out by the new user is valid 226, an appropriate SID is generated 228. If the registration is not valid, access is again denied 222.

An SID for an authorized user is appended ("tagged") 230 to the original URL directed to a controlled page on the content server. The authentication server then transmits a REDIRECT response 232 based on the tagged URL to the client browser 100. The modified URL, such as "http://content.com/[SID]/report" is automatically forwarded to the content server 120.

Figure 3, illustrates a typical client-server exchange involving the access control and monitoring method of the present invention. In Step 1, the client 50 running a browser transmits a GET request through a network for an uncontrolled page (UCP). For example, the user may request an advertisement page by transmitting a URL "http://content.com/advertisement", where "content.com" is the server name and "advertisement" is the uncontrolled page name. In Step 2, the content server 52 processes the GET request and transmits the requested page, "advertisement". The content server also logs the GET request in the transaction database 56 by recording the URL, the client IP address, and the current time.

In Step 3, the user on the client machine may elect to traverse a link in the advertisement page directed to a controlled page (CP). For example, the advertisement page may contain a link to a controlled page called "report". Selecting this link causes the client browser 50 to forward a GET request through a URL which is associated with the report file "http://content.com/report". The content server 52 determines that the request is to a controlled page and that the URL does not contain an SID. In Step 4, the content server transmits a REDIRECT response to the client, and, in Step 5, the browser automatically sends the REDIRECT URL to the authentication server 54. The REDIRECT URL sent to the authentication server may contain the following string:

10 "http://auth.com/authenticate?domain=[domain]&URL=http://
content.com/report"

The authentication server processes the REDIRECT and determines whether user credentials (CRED) are needed for authorization. In Step 6, the authentication server transmits a "CHALLENGE" response to the client. As previously described, typical credentials consist of user name and password. An authorization header based on the credential information is then forwarded by the client browser to the authentication server. For example, a GET URL having such an authorization header is: "http://autho.com/authenticate?domain=[domain]&URL=http://content.com/report and the authorization header may be:"AUTHORIZE:[authorization]". The authentication server processes the GET request by checking the Account Database 58. If a valid account exists for the user, an SID is issued which authorizes access to the controlled page "report" and all the other pages within the domain.

As previously described, the preferred SID comprises a compact ASCII string that encodes a user identifier, the current domain, a key identifier, an expiration time, the client IP address, and an unforgeable digital signature. In Step 8, the authentication

server redirects the client to the tagged URL, "http://content.com/[SID]/report", to the client. In Step 9, the tagged URL is automatically forwarded by the browser as a GET request to the content server. The content server logs the GET request in the Transaction database 56 by recording the tagged URL, the client IP address, and the
5 current time. In Step 10, the content server, upon validating the SID, transmits the requested controlled page "report" for display on the client browser.

According to one aspect of the present invention, the content server periodically evaluates the record contained in the transaction log 56 to determine the frequency and duration of accesses to the associated content server. The server counts requests to
10 particular pages exclusive of repeated requests from a common client in order to determine the merits of the information on different pages for ratings purposes. By excluding repeated calls, the system avoids distortions by users attempting to "stuff the ballot box."

In one embodiment, the time intervals between repeated requests by a common
15 client are measured to exclude those requests falling within a defined period of time.

Sub P2
Additionally, the server may, at any given time, track access history within a client-server session. Such a history profile informs the service provider about link transversal frequencies and link paths followed by users. This profile is produced by filtering transaction logs from one or more servers to select only transactions involving
20 a particular user ID (UID). Two subsequent entries, A and B, corresponding to requests from a given user in these logs represent a link traversal from document A to document B made by the user in question. This information may be used to identify the most popular links to a specific page and to suggest where to insert new links to provide more direct access. In another embodiment, the access history is evaluated to determine
25 traversed links leading to a purchase of a product made within commercial pages. This information may be used, for example, to charge for advertising based on the number of

link traversals from an advertising page to a product page or based on the count of purchases resulting from a path including the advertisement. In this embodiment, the server can gauge the effectiveness of advertising by measuring the number of sales that resulted from a particular page, link, or path of links. The system can be configured to charge the merchant for an advertising page based on the number of sales that resulted from that page.

According to another aspect of the present invention, a secondary server, such as the authentication server 200 in Figure 2B, may access a prearranged user profile from the account database 216 and include information based on such a profile in the user identifier field of the SID. In a preferred embodiment, the content server may use such an SID to customize user requested pages to include personalized content based on the user identifier field of the SID.

In another aspect of the invention, the user may gain access to domain of servers containing journals or publications through a subscription. In such a situation, the user may purchase the subscription in advance to gain access to on-line documents through the Internet. The user gains access to a subscribed document over the Internet through the authorization procedure as described above where an authorization indicator is preferably embedded in a session identifier. In another embodiment, rather than relying on a prepaid subscription, a user may be charged and billed each time he or she accesses a particular document through the Internet. In that case, authorization may not be required so long as the user is fully identified in order to be charged for the service. The user identification is most appropriately embedded in the session identifier described above.

In another aspect of the invention, facilities are provided to allow users to utilize conventional telephone numbers or other identifiers to access merchant services. These merchant services can optionally be protected using SIDs. In a preferred embodiment,

as shown in Figure 6, a Web browser client 601 provides a "dial" command to accept a telephone number from a user, as by clicking on a "dial" icon and inputting the telephone number through the keyboard. The browser then constructs a URL of the form "http://directory.net/NUMBER", where NUMBER is the telephone number or
5 other identifier specified by the user. The browser then performs a GET of the document specified by this URL, and contacts directory server 602, sending the NUMBER requested in Message 1.

In another embodiment, implemented with a conventional browser, client 601
10 uses a form page provided by directory server 601 that prompts for a telephone number or other identifier in place of a "dial" command, and Message 1 is a POST message to a URL specified by this form page.

Once NUMBER is received by directory server 601, the directory server uses
database 604 to translate the NUMBER to a target URL that describes the merchant
server and document that implements the service corresponding to NUMBER. This
15 translation can ignore the punctuation of the number, therefore embedded parenthesis or dashes are not significant.

In another embodiment an identifier other than a number may be provided. For
example, a user may enter a company name or product name without exact spelling. In
such a case a "soundex" or other phonetic mapping can be used to permit words that
20 sound alike to map to the same target URL. Multiple identifiers can also be used, such
as a telephone number in conjunction with a product name or extension.

In Message 2, Directory server 602 sends a REDIRECT to client 601, specifying
the target URL for NUMBER as computed from database 604. The client browser 601
then automatically sends Message 3 to GET the contents of this URL. Merchant server
25 603 returns this information in Message 4. The server 602 might have returned a Web

002140-52237560

5 The Target URL contained in Message 3 can be an ordinary URL to an uncontrolled page, or it can be a URL that describes a controlled page. If the Target URL describes a controlled page then authentication is performed as previously described. The Target URL can also describe a URL that includes an SID that provides a preauthorized means of accessing a controlled page.

15 In the approach a single merchant server can provide multiple services that correspond to different external "telephone numbers" or other identifiers. For example, if users dial the "flight arrival" number they could be directed to the URL for the arrival page, while, if they dial the "reservations" number, they would be directed to the URL for the reservations page. A "priority gold" number could be directed to a controlled
20 page URL that would first authenticate the user as belonging to the gold users group, and then would provide access to the "priority gold" page. An unpublished "ambassador" number could be directed to a tagged URL that permits access to the "priority gold" page without user authentication.

This invention has particular application to network sales systems such as presented in U.S. Patent Application Serial No. 08/328,133, filed October 24, 1994, by Payne et al. which is incorporated herein by reference.

5 EQUIVALENTS

Those skilled in the art will know, or be able to ascertain using no more than routine experimentation, many equivalents to the specific embodiments or the invention described herein. These and all other equivalents are intended to be encompassed by the following claims.

00210-5E28150